# A Note on Privacy in Constant Function Market Makers

Guillermo Angeris
angeris@stanford.edu

Alex Evans
alex@placeholder.vc

Tarun Chitra
tarun@gauntlet.network

February 2021

**Abstract**

Constant function market makers (CFMMs) such as Uniswap, Balancer, Curve, and mStable, among many others, make up some of the largest decentralized exchanges on Ethereum and other blockchains. Because all transactions are public in current implementations, a natural next question is if there exist similar decentralized exchanges which are privacy-preserving; *i.e.*, if a transaction's quantities are hidden from the public view, then an adversary cannot correctly reconstruct the traded quantities from other public information. In this note, we show that privacy is impossible with the usual implementations of CFMMs under most reasonable models of an adversary and provide some mitigating strategies.

# Introduction

Decentralized exchanges (DEXs) have experienced rapid growth in liquidity and trading volume over the last year. Much of this growth can be attributed to the rise of constant function market makers (CFMMs) that allow for computationally cheap on-chain trading [AC20]. This growth has, in turn, motivated attempts to improve existing mechanisms for decentralized exchange. For example, current DEX designs do not support private trading, as the full details of each trade that users make can be directly attributed to their on-chain identities. In addition to other challenges, the lack of privacy also makes it easier for third parties to front-run a user's trades [DGK+20, TCS21]. Another such problem is the ability for attackers to deanonymize agents by doing basic statistical analyses of public trades performed on DEXs [GGS20, Cha20]. In contrast, centralized brokers and exchanges preserve user privacy, but agents are required to trust that the exchange won't leak sensitive trade data. A natural question to ask is whether or not popular decentralized exchanges such as Uniswap [AKC+19, ZR] can be adapted to preserve privacy.

The advent of smart contract systems that utilize zero-knowledge proof systems, such as Zexe [BCG+20], suggest that it should be possible to privately execute CFMM transactions. Indeed, a number of proposed protocols such as SecretSwap [Pow21] and Manta [CXZ19], propose potentially privacy-preserving modifications to Uniswap via the use of either trusted

hardware or cryptographic improvements. However, it has been informally and heuristically noted that 'black-box' applications of privacy-preserving technology to Uniswap are unlikely to preserve privacy [Whi20] as the timing of a trade implicitly leaks identity within Uniswap and other constant function market makers (CFMMs), and can be used to reconstruct the trade.

In this paper, we formalize this intuition and prove that CFMMs are generically unable to preserve privacy under even relatively weak adversaries. We construct a model where knowledge of a CFMM trading function or 'invariant,' such as Uniswap's famous $xy = k$ model, combined with observations of the time-ordering of trades allows an attacker to recover the traded quantities, provided the agent is able to interact with the CFMM contract in a meaningful way. One of the main benefits of the convexity of a CFMM is that it makes the arbitrage problem between exchanges easy [AKC+19]. Our results illustrate a downside to this: convexity allows an adversary to uniquely recover the traded quantities, assuming that neither the reserve values nor the traded amounts are known.

**Summary.** We give a very basic introduction to CFMMs and describe the attack in §1. The attack depends on a uniqueness result for a certain system of nonlinear equations, which we show in §1.3 by some basic tools of convex analysis, for a relatively general family of trading functions. In §2 we provide some basic extensions of the proof to more general CFMMs and slightly weaker attacker models. Finally, we provide a number of potential mitigation mechanisms that protocol designers can use to improve user privacy in §3. While our results are negative, they illustrate that more complex economic mechanisms are needed to preserve privacy than one might initially assume. We hope these results can be used to guide future private decentralized exchange design.

# 1 Impossibility of privacy

We will show that constant function market makers cannot be private in their usual implementations, under most reasonable models of adversaries. For simplicity, we assume no fees, but discuss an extension to the case with fees in §2.

## 1.1 Constant function market makers

We provide only topically relevant definitions of CFMMs in this short note and refer the reader to [AC20] for a thorough introduction to both the definitions and many of the tools used throughout this paper. The notation here differs slightly for simplicity of the presentation, but is equivalent to the sufficient condition for path independence in [AC20, §2.3.2] in the case of two assets.

**Definition.** A *constant function market maker*, or CFMM, is an automated market maker defined by its *reserve quantities* $R \in \mathbf{R}^n_+$ and a *trading function* $\psi : \mathbf{R}^n_+ \to \mathbf{R}$. The behavior of CFMMs is very simple: an agent proposes some trade $\Delta \in \mathbf{R}^n$ where $\Delta_i$ is a positive

quantity if $\Delta_i$ of coin $i$ is given to the CFMM while it denotes a negative quantity if it is taken. The CFMM then checks if the trade satisfies

$$\psi(R + \Delta) = \psi(R),$$

*i.e.*, if the trade function $\psi$, depending on the reserves, does not change in value after the proposed trade. If so, the trade is accepted and the CFMM pays out $-\Delta$ of the traded asset from its reserves $R$, leading to the reserve values being updated as $R \leftarrow R + \Delta$. If not, the trade is rejected and the CFMM does not change its state. Additionally, trades for which $R + \Delta \not\geq 0$ are always rejected since they cannot be fulfilled with the current reserves.

**Assumptions.** We will, in general, assume that the trading function $\psi$ is a strictly concave, increasing function, which holds for essentially all CFMMs barring some special cases such as mStable (or constant sum market makers). This is true for any CFMM whose reachable set [AC20, §2.3] is a strictly convex set, for all reserves $R$. For example, in the case of Uniswap, or constant product markets, $\psi(R) = R_1 R_2$, which is neither concave nor convex, but it can be equivalently written as $\psi(R) = \sqrt{R_1 R_2}$, which is strictly concave, increasing whenever $R > 0$. (The notion of equivalence used here is that of [AC20, §2.1], which we will not discuss further in this note.) We discuss extensions which include trading functions that are not strictly concave in §2, while we may generally assume that the function $\psi$ is nondecreasing without loss of generality [AC20, §A.1].

**Reported price.** As shown in [AC20, §2.4], we have that the marginal price, $c \in \mathbf{R}^n_+$ of a fee-less CFMM with reserves $R$ is given by

$$\nabla\psi(R) = \lambda c,$$

where $\lambda \geq 0$ is a nonnegative scalar multiplier.

## 1.2 Adversary definition and attack

We assume a very simple, but very general, model of an adversary. In our case, the adversary, who we will call Eve, attempts to discover the quantity traded by an agent, called Alice. In our model, Eve is unable to see the exact quantities Alice used to trade with the CFMM, but knows when Alice's transaction took place. Eve's only ability is to interact with the CFMM in a state before Alice's transaction and after the transaction.

**Action space.** In our case, we will assume that Eve is able to query the marginal price of the CFMM, at the current reserves, and whether a given trade $\Delta$ is valid. (We assume she has access to at least one nonzero valid trade.) We will also assume, as is generally the case, that Eve knows when Alice's transaction took place and can query the CFMM in its state before and after the transaction. This assumption could be broken, *e.g.*, if there exist fully-private protocols in which no agent can know the transaction times of any other

agent, but such protocols have not yet made it into production.[1] We note that this attacker model, where Eve knows only the transaction time, is relatively different than standard attacker models in the blockchain setting [GKL15, BGK+19] and might therefore be useful to consider in their own right.

**Attack description.** The attack will make repeated use of the following 'atom': Eve is always able to reconstruct the reserve amounts given (a) the marginal price at the current reserves and (b) a single nonzero feasible trade, by solving a basic nonlinear system of equations. Using this, it is then enough to simply compute the reserve amounts before and after Alice's trade to recover the traded amounts. More explicitly, the sequence of the attack is as follows:

1. Eve queries the marginal price of the CFMM at the current reserves, to get some vector $c$ and then queries any valid nonzero trade $\Delta \neq 0$.

2. Using this information and the known functional form of $\psi$, Eve can recover the reserves $R \in \mathbf{R}^n_{++}$ by finding a solution to the following nonlinear system of equations in $R$:

$$\nabla\psi(R) = \lambda c, \quad \psi(R + \Delta) = \psi(R). \tag{1}$$

   The fact that this system has a unique solution (*i.e.*, the true reserves $R$) is a slightly technical point which we will discuss later in this section.

3. Letting Alice's trade be $\Delta_a$ the new reserves are $R_a = R + \Delta_a$, which are not known to Eve, but the CFMM can now be queried in this new state.

4. Eve then queries the contract again to get a new marginal price $c'$ and then queries any nonzero trade $\Delta'$. She again solves the corresponding system of equations (1) to find the new reserves $R_a$:

$$\nabla\psi(R_a) = \lambda c', \quad \psi(R_a + \Delta') = \psi(R_a),$$

   to receive $R_a \in \mathbf{R}^n_{++}$.

5. Eve then computes $R_a - R = \Delta_a$ to receive Alice's traded values.

From here, it is clear that Eve can always exactly compute the traded amounts from Alice, even if Eve is only given access to very basic quantities. The only thing that remains to be shown is the uniqueness of the solution $R$ and $R_a$ to the system of equations. (Existence is guaranteed since the true reserve values $R$ and $R_a$ satisfy the equations, by definition.) There are, of course, many ways in which Eve can compute a solution to (1) given her known data $c$ and $\Delta$. For example, using a Newton-type method will likely yield very good practical results for general $\psi$, but the results are much simpler in some important special cases.

[1] Cryptographic primitives such as Verifiable Delay Functions (VDFs) [BBBF18] can provide such transaction randomization, but have yet to be used in production networks, let alone those with high-transaction rates.

**Reserve discovery in Uniswap.** In the case where $\psi(R)$ is a constant product market maker such as Uniswap, *i.e.*, when $R \in \mathbf{R}^2_{++}$ and

$$\psi(R) = \sqrt{R_1 R_2},$$

then (1) reduces to a linear system of equations in $R_1$ and $R_2$. In particular, we have:

$$\nabla\psi(R) = \frac{\sqrt{R_1 R_2}}{2}\left(\frac{1}{R_1}, \frac{1}{R_2}\right) = \lambda(c_1, c_2),$$

so

$$\frac{1}{2}\sqrt{\frac{R_2}{R_1}} = \lambda c_1, \qquad \frac{1}{2}\sqrt{\frac{R_1}{R_2}} = \lambda c_2. \tag{2}$$

Multiplying both sides of each equation gives

$$\frac{1}{4} = \lambda^2 c_1 c_2,$$

or that $\lambda = (2\sqrt{c_1 c_2})^{-1}$, since $\lambda \geq 0$. Plugging this value back into (2), we find that

$$\frac{R_2}{R_1} = \frac{c_1}{c_2},$$

or that $c_1 R_1 = c_2 R_2$. Finally, let $\Delta \in \mathbf{R}^2$ be any feasible trade, then

$$\psi(R + \Delta) = \sqrt{(R_1 + \Delta_1)(R_2 + \Delta_2)} = \psi(R) = \sqrt{R_1 R_2},$$

which easily simplifies to

$$\Delta_2 R_1 + \Delta_1 R_2 + \Delta_1 \Delta_2 = 0.$$

We can then easily recover $R_1$ and $R_2$, given $c$ and $\Delta$ by solving the following system of linear equations:

$$\begin{aligned} c_1 R_1 - c_2 R_2 &= 0 \\ \Delta_2 R_1 + \Delta_1 R_2 &= -\Delta_1 \Delta_2. \end{aligned}$$

This system has a unique solution $(R_1, R_2)$ provided $c_1\Delta_1 + c_2\Delta_2 \neq 0$, which can be shown to hold for all feasible trades $\Delta \neq 0$. (We provide a much more general proof, which includes this as a special case, in §1.3.)

The fact that Uniswap's reserves can be recovered using only the marginal price $c$ and a nonzero feasible trade $\Delta$ has at least one simple, direct proof, which does not make use of (1), but we provide this special case as an easily-verifiable example of the more general attack. In fact, in the special case of constant product markets, an adversary only requires the existence of any two nonzero, distinct feasible trades in order to correctly reconstruct the reserves at some given point in time. We encourage the reader to try this specific problem as an exercise, and show that this method extends generally to other CFMMs in §2.

## 1.3 Uniqueness of solution

We will show that the solution of (1) is unique in the case that $\psi$ is an increasing, nonnegative, strictly concave function that is 1-homogeneous; *i.e.*, when

$$\psi(kR) = k\psi(R)$$

for any $k \geq 0$. (This includes constant product and constant mean markets, such as Uniswap and Balancer, as special cases.) We suspect that uniqueness of the solution can be shown in the more general setting where the function is not 1-homogeneous but leave this for future work.

**Reserves at fixed price.** We will define the set of reserves consistent with the first constraint as:

$$Q(c) = \{R > 0 \mid \nabla\psi(R) = \lambda c, \text{ for some } \lambda \geq 0\}.$$

In other words, $Q(c)$ is the set of reserves which are consistent with the marginal price of $c$.

Because $\psi$ is a strictly concave, 1-homogeneous function, we will show that this set is a ray, *i.e.*, it can be written as

$$Q(c) = \{kR^0 \mid k > 0\},$$

for any $R^0 \geq 0$ with $\psi(R^0) > 0$, satisfying $\nabla\psi(R^0) = \lambda^0 c$ for some $\lambda^0 \geq 0$. Note that inclusion, $\{kR^0 \mid k > 0\} \subseteq Q(c)$, follows immediately from the fact that $\psi$ is 1-homogeneous and $R^0 \in Q(c)$. On the other hand, showing that $Q(c) \subseteq \{kR^0 \mid k > 0\}$ is slightly tricker.

To do this, start with any $R \in Q(c)$ and consider the $\alpha$-superlevel set of $\psi$, given by:

$$S(\alpha) = \{R \mid \psi(R) \geq \alpha\},$$

which is a strictly convex set since $\psi$ is a strictly concave function. Additionally, we will make use of the fact that, for any $k > 0$,

$$kS(\alpha) = S(k\alpha),$$

by the homogeneity of $\psi$. (Here $kS(\alpha)$ denotes elementwise set multiplication.)

Given this definition, we know that $c$ is a supporting hyperplane of the set $S(\psi(R^0))$ at the point $R^0$ and of the set $S(\psi(R))$ at the point $R$. (This follows immediately from the first-order conditions for convexity applied to the function $\psi$ along with the definition of $Q(c)$.) Now, because $\psi(R^0) > 0$ and $\psi(R) > 0$, there exists some $k > 0$ such that $k\psi(R^0) = \psi(R)$. Additionally, we have, by homogeneity,

$$kS(\psi(R^0)) = S(k\psi(R^0)) = S(\psi(kR^0)) = S(\psi(R)).$$

But, since $c$ is a supporting hyperplane for $S(\psi(R^0))$ at $R^0$, it is a supporting hyperplane of $kS(\psi(R^0))$ and therefore of $S(\psi(kR^0)) = S(\psi(R))$ at $kR^0$ and $R$. By strict convexity, every supporting hyperplane of a set will map to a unique point on the boundary, so we must have that, in fact, $kR^0 = R$, so $R$ lies on the ray generated by $R^0$, as required.

**Reserves consistent with a trade.** Now we have to show that the intersection between the set $Q(c)$ and the set

$$U(\Delta) = \{R > 0 \mid \psi(R + \Delta) = \psi(R)\},$$

is a singleton; *i.e.*, that there is a unique solution to the nonlinear system given in (1). We can interpret the set $U(\Delta)$ as the set of reserves for which a given trade $\Delta$ is feasible. Note that any solution to (1), for given $\Delta$ and $c$, is, by definition, going to be in the intersection of $U(\Delta) \cap Q(c)$. Because the true reserves $R$ satisfy both equations, it is clear that $U(\Delta) \cap Q(c)$ is nonempty; our goal now is to show that the intersection contains exactly one element, $R$, and therefore that Eve can correctly recover the reserves by finding a solution to (1).

To show that this intersection is a singleton, it suffices to show that

$$\psi(kR + \Delta) = \psi(kR)$$

has the unique solution $k = 1$, because $R \in Q(c)$ and therefore every element of $Q(c)$ is of the form $kR$ by the previous argument. We will divide this problem into the cases where $k > 1$ and $k < 1$. First, assume that $k > 1$, then, by definition of $R$,

$$\psi(R + \Delta) = \psi(R),$$

but since $\psi$ is strictly concave, then, for any $0 < \eta < 1$:

$$\psi(R + \eta\Delta) = \psi(\eta(R + \Delta) + (1 - \eta)R) > \eta\psi(R + \Delta) + (1 - \eta)\psi(R) = \psi(R).$$

Setting $\eta = 1/k$ we have that $0 < \eta < 1$, so

$$\psi(R + (1/k)\Delta) > \psi(R),$$

or, multiplying on both sides by $k$ and using the homogeneity of $\psi$:

$$\psi(kR + \Delta) > \psi(kR),$$

so $k > 1$ cannot be a solution. To show the $0 < k < 1$ case, we will show the contrapositive: if $0 < k < 1$ and

$$\psi(kR + \Delta) \geq \psi(kR),$$

then $\psi(R + \Delta) \neq \psi(R)$. This follows from a nearly identical proof as the above: we have that

$$\psi(kR + k\Delta) > k\psi(kR + \Delta) + (1 - k)\psi(kR) \geq \psi(kR),$$

where the first inequality follows from the strict concavity of $\psi$, while the second follows by assumption. Then we immediately have:

$$\psi(R + \Delta) = \frac{1}{k}\psi(kR + k\Delta) > \frac{1}{k}\psi(kR) = \psi(R),$$

so $\psi(R + \Delta) \neq \psi(R)$ as required. The contrapositive then implies that, if $R$ is a solution, we must have that

$$\psi(kR + \Delta) < \psi(kR).$$

for $0 < k < 1$. Combining both statements gives that

$$\psi(kR + \Delta) = \psi(kR)$$

if, and only if, $k = 1$.

**Discussion.** The proof essentially makes use of two important 'tricks,' which might be generalizable to the case where the function $\psi$ is not homogeneous of any nonzero degree.

The first is that, because the $\alpha$-superlevel set of $\psi$ is strictly convex, then any supporting hyperplane maps to a unique reserve value $R$ (depending, implicitly, on $\alpha$). In some sense, this provides a way of 'identifying' reserves, at some fixed, but potentially unknown, liquidity, with a marginal price $c$. In our case, we used the homogeneity to prove an explicit form for the set $Q(c)$, but this is likely unnecessary and $Q(c)$ likely satisfies some similarly useful property without requiring the homogeneity of $\psi$. (We note that Curve is a counterexample to the plausible conjecture that $Q(c)$ is always a ray for all strictly convex trading functions, since scaling the reserves of Curve by any nonzero constant will change the marginal price with respect to any numéraire whenever $\alpha, \beta > 0$; *cf.*, [AC20, §2.4].)

The second is that, by the monotonicity and strict concavity of $\psi$, a given trade $\Delta$ feasible for some reserves $R$ with marginal price $c$ will either be too expensive when the CFMM has more liquidity than $R$ (*i.e.*, there is a strictly better trade $\Delta'$ for the trader that is feasible) or infeasible when there is less liquidity than $R$. In the case where $\psi$ is homogeneous, 'more' and 'less' liquidity, at fixed marginal price $c$, are very easy to identify, since all possible reserves lie along a ray and are totally ordered, but a more general construction will require some care.

# 2 Extensions

There are a number of basic extensions which are available to this attack and for which the proof still holds with either slight or no modifications.

**Nonzero fees.** In the case that the function has nonzero fees, *i.e.*, if the CFMM must instead satisfy

$$\psi(R + \gamma\Delta_+ - \Delta_-) = \psi(R),$$

where $\Delta_+$ is the vector whose nonzero entries are the nonnegative entries of $\Delta$ (with all other entries equal to zero) and similarly for $\Delta_-$, except with the nonpositive entries of $\Delta$. In this case, the 'feasibility' condition is changed slightly for a trade, but the proof of uniqueness remains otherwise identical.

**General homogeneity.** Although we assume 1-homogeneity for a slightly cleaner exposition, the proof is nearly identical if 1-homogeneity of $\psi$ is replaced with $p$-homogeneity of $\psi$; *i.e.*, for $\lambda \geq 0$

$$\psi(\lambda R) = \lambda^p \psi(R),$$

with $p \neq 0$. The case of $p = 0$ is unlikely to be useful since it would imply that the CFMM is not sensitive with respect to scaling of the reserves (*i.e.*, it is liquidity-insensitive) which would imply that liquidity provision does not change the dynamics of the CFMM. We expect similar problems in liquidity-insensitive CFMMs as those of the classical AMMs [OS11].

**Unknown marginal price.** In the case that the marginal price is unknown; *e.g.*, it cannot be accessed directly, it is not difficult for Eve to compute an arbitrarily-good approximation by performing $n$ queries. In particular, let $\Delta^i \in \mathbf{R}^n$ for $i = 1, \ldots, n$ be any $n$ feasible trades, then we know that, by the concavity of $\psi$:

$$\psi(R + \Delta^i) < \psi(R) + \lambda c^T \Delta^i, \qquad i = 1, \ldots, 2n,$$

where $\lambda c = \nabla\psi(R)$ with $\lambda > 0$ fixed. But, since $\Delta^i$ is a feasible trade, we have that $\psi(R) = \psi(R + \Delta^i)$, so $c$ must satisfy

$$\begin{aligned} c^T \Delta^i &\geq \varepsilon \mathbf{1}, \quad i = 1, \ldots, n \\ c &\geq 0, \end{aligned}$$

where $\mathbf{1}$ is the all-ones vector and $\varepsilon > 0$ satisfies

$$\varepsilon \leq \min_i(\psi(R) - \psi(R + \Delta^i)).$$

We note that this is not essential, since the equations can be scaled by $1/\varepsilon$ to recover $c$ up to a constant multiple, so it suffices to solve the following system of inequalities:

$$\begin{aligned} c^T \Delta^i &\geq \mathbf{1}, \quad i = 1, \ldots, n \\ c &\geq 0. \end{aligned} \tag{3}$$

This sets up a system of $2n$ inequalities in $n$ variables, which has a unique solution, up to a nonnegative constant multiple, provided some basic conditions on the trades $\Delta^i$. In particular, uniqueness of a solution can be shown since $\psi$ is strictly concave, under some mild conditions on the queries. The attack then proceeds as previously stated, replacing querying the marginal price with querying $n$ feasible trades and computing the marginal price by solving the system of inequalities (3).

**Non-strict concavity.** The more general case where $\psi$ is not strictly concave is rather more difficult because it implies that there exists some reserve quantities which may map to the same marginal price and the uniqueness proof need not hold. In some special cases, such as constant sum market makers (*e.g.*, mStable), there are simple mitigating strategies: since the marginal price is fixed everywhere, Eve can make proportionally larger trades until a trade becomes infeasible. She can then simply perform a binary search to get $\varepsilon$-close to the true reserve quantities in $\log(1/\varepsilon)$ queries.

More general CFMMs whose concave functions are not strictly concave are slightly more difficult, but follow by a similar argument. In that case, given a fixed marginal price $c$, the set of reserves consistent with $c$, defined as $Q(c)$ above, is a set whose affine hull might have dimension greater than 1, but is a convex set. Because these reserves all imply the same marginal price, Eve can similarly query progressively larger trades at the fixed price, until these trades become infeasible. Using the fact that a point in a convex set can be

identified with $n+1$ distances to known points on the set (under mild conditions), Eve can then identify a unique reserve $R \in Q(c)$ satisfying all of the above constraints.

Making the uniqueness argument rigorous is rather more involved, even in the case where $\psi$ is 1-homogeneous, and we do not make it here as most CFMMs are strictly concave in practice.

# 3 Mitigating strategies

There are a number of useful conditions for which the proof fails and such conditions might lead to modifications of CFMMs in order to guarantee privacy. We discuss some basic examples of what these modifications might look like, but do not prove (nor guarantee) that these modifications are sufficient for privacy.

**Randomness in price.** One immediate example is that Eve requires knowledge of $c$ in order to reconstruct the reserves. Here, it is possible that the CFMM could add some amount of randomness (in a similar vein to differential privacy; see, *e.g.*, [BD14, DR$^+$14]) to the price in order to prevent Eve from correctly reconstructing the reserve values. Note that this randomness must be fixed at each block since it would otherwise be easy to approximate the true price by averaging multiple queries. Additionally, we note that this construction can quickly become complicated since the price, with randomness, needs to be consistent with the feasible trades and cannot 'leak' too much information after the trades are completed. There is also the more general problem that such differences in price, due to the added randomness, might be exploited by arbitrageurs, causing additional losses for the liquidity providers.

**Batching orders.** Another possibility is that the CFMM automatically *batches* orders; *i.e.*, the CFMM waits until several trades $\Delta$ are accepted and updates its reserves only after all trades have been executed, taking the trades and paying the output all at once. This leads to two potential difficulties. The first is that the CFMM needs to ensure that all of the orders, excluding Alice's, are not under Eve's control. (This could be guaranteed, *e.g.*, if there is always enough trading volume that could make this attack prohibitively expensive for Eve.) The second is that the CFMM needs to wait until a batch of orders is specified before executing the trades. This delay can lead to bad user-facing performance and may threaten the solvency of the system in extreme cases with large price fluctuations.

**Further thoughts.** In general, there will likely always be a tradeoff between the cost users (*i.e.*, liquidity providers or traders) are willing to pay in order to achieve privacy. Such a cost can be a direct cost, such as higher prices for a fixed trade, or an indirect cost, such as higher risk of systemic failure, resulting in higher price volatility within such systems. At the moment, it is not clear what price a given proportion of users are willing to pay in order to ensure privacy. This leads to further question on where we should allocate the

(finite) developer time available for building such systems in order to ensure the best user experience, while guaranteeing that the systems are safe in useful ways.

# 4 Conclusion

We have shown that privacy in CFMMs requires more than simply obscuring reserve and trade quantities from possible attackers. In particular, an attacker can always recover the reserve quantities of a CFMM given only information that is required for agents to be able to interact with the system in a meaningful way: *i.e.*, given some amount of coin Alice adds to the CFMM, how much output can she receive? This attack highlights the difficulty of achieving privacy in DeFi even under relatively weak adversarial models. We also discussed possible ways of preventing attackers from knowing the traded quantities, but note that these are either difficult to achieve in practice or suffer from a degraded user experience. In general, we suspect that there exist reasonable variations of CFMMs which are privacy-preserving, but their implementation is likely to be neither obvious nor frictionless for the end user.

This note leaves open two major research questions, in order of increasing importance. First, does the attack outlined here work for all strictly convex CFMMs, not just those that are 1-homogeneous? We suspect this is the case, but have not been able to give a reasonable proof except in other special cases. And, second, what does a privacy-preserving CFMM look like, and what privacy guarantees can it make? This question is likely much harder, but also far more important. We suspect that any reasonable progress made towards answering it is likely to be very useful in both the theory and practice.

# Acknowledgements

The authors would like to thank Assimakis Kattis for useful discussions regarding mitigating strategies and Anna Rose for inspiring this note.

# References

[AC20] Guillermo Angeris and Tarun Chitra. Improved Price Oracles: Constant Function Market Makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pages 80–91, New York NY USA, October 2020. ACM.

[AKC+19] Guillermo Angeris, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra. An analysis of Uniswap markets. *Cryptoeconomic Systems*, 2019.

[BBBF18] Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In *Annual international cryptology conference*, pages 757–788. Springer, 2018.

[BCG+20] Sean Bowe, Alessandro Chiesa, Matthew Green, Ian Miers, Pratyush Mishra, and Howard Wu. Zexe: Enabling decentralized private computation. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 947–964. IEEE, 2020.

[BD14] Rina Foygel Barber and John C Duchi. Privacy and statistical risk: Formalisms and minimax bounds. *arXiv preprint arXiv:1412.4451*, 2014.

[BGK+19] Christian Badertscher, Peter Gazi, Aggelos Kiayias, Alexander Russell, and Vassilis Zikas. Ouroboros chronos: Permissionless clock synchronization via proof-of-stake. *IACR Cryptol. ePrint Arch.*, 2019:838, 2019.

[Cha20] Chainalysis. The 2020 state of crypto crime. 2020.

[CXZ19] Shumo Chu, Qiudong Xia, and Zhenfei Zhang. Manta: Privacy preserving decentralized exchange. 2019.

[DGK+20] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 910–927. IEEE, 2020.

[DR+14] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.

[GGS20] Daniel Goldsmith, Kim Grauer, and Yonah Shmalo. Analyzing hack subnetworks in the bitcoin transaction graph. *Applied Network Science*, 5(1):1–20, 2020.

[GKL15] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In *Annual international conference on the theory and applications of cryptographic techniques*, pages 281–310. Springer, 2015.

[OS11] Abraham Othman and Tuomas Sandholm. Liquidity-Sensitive Automated Market Makers via Homogeneous Risk Measures. In Ning Chen, Edith Elkind, and Elias Koutsoupias, editors, *Internet and Network Economics*, volume 7090, pages 314–325. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.

[Pow21] Benjamin Powers. Secretswap is the secret network's answer to defi privacy, Feb 2021.

[TCS21] Christof Ferreira Torres, Ramiro Camino, and Radu State. Frontrunner jones and the raiders of the dark forest: An empirical study of frontrunning on the ethereum blockchain. *arXiv preprint arXiv:2102.03347*, 2021.

[Whi20] Barry WhiteHat. Why you can't build a private Uniswap with ZKPs. `https://ethresear.ch/t/why-you-cant-build-a-private-uniswap-with-zkps/7754`, Jul 2020.

[ZR] Noah Zinsmeister and Dan Robinson. Uniswap v2 Core. page 10.